

Spin-SILC: CMB polarisation component separation with spin wavelets

Keir K. Rogers,^{1*} Hiranya V. Peiris,¹ Boris Leistedt,^{2,1} Jason D. McEwen³
and Andrew Pontzen¹

¹*Department of Physics & Astronomy, University College London, Gower Street, London WC1E 6BT, UK*

²*Department of Physics, New York University, 4 Washington Place, New York, NY 10003, USA*

³*Mullard Space Science Laboratory, University College London, Surrey RH5 6NT, UK*

23 August 2016

ABSTRACT

We present Spin-SILC, a new foreground component separation method that accurately extracts the cosmic microwave background (CMB) polarisation E and B modes from raw multifrequency Stokes Q and U measurements of the microwave sky. Spin-SILC is an internal linear combination method that uses spin wavelets to analyse the spin-2 polarisation signal $P = Q + iU$. The wavelets are additionally directional (non-axisymmetric). This allows different morphologies of signals to be separated and therefore the cleaning algorithm is localised using an additional domain of information. The advantage of spin wavelets over standard scalar wavelets is to simultaneously and self-consistently probe scales and directions in the polarisation signal $P = Q + iU$ and in the underlying E and B modes, therefore providing the ability to perform component separation and E - B decomposition concurrently for the first time. We test Spin-SILC on full-mission *Planck* simulations and data and show the capacity to correctly recover the underlying cosmological E and B modes. We also demonstrate a strong consistency of our CMB maps with those derived from existing component separation methods. Spin-SILC can be combined with the pseudo- and pure E - B spin wavelet estimators presented in a companion paper to reliably extract the cosmological signal in the presence of complicated sky cuts and noise. Therefore, it will provide a computationally-efficient method to accurately extract the CMB E and B modes for future polarisation experiments.

Key words: cosmic background radiation – methods: data analysis

1 INTRODUCTION

The polarisation of the cosmic microwave background (CMB) is a powerful cosmological observable, providing deep insights into the physics of the early universe. The decomposition of the linear polarisation into curl-free (E mode) and divergence-free (B mode) components allows the detection of tensor perturbations to the metric. Specifically, a non-zero BB power spectrum on degree scales would support the existence of a stochastic background of gravitational waves predicted by inflationary theory (Kamionkowski et al. 1997; Seljak & Zaldarriaga 1997). Accurate measurement of B mode polarisation on arcminute scales also gives strong constraints on the neutrino sector via the weak gravitational lensing of CMB E modes (Zaldarriaga & Seljak 1998). There are numerous existing and planned ground-based, balloon-borne and satellite experiments designed to precisely measure CMB polarisation (see e.g., Errard et al. 2015 for a recent forecast on the cosmological constraining power of current and upcoming missions). However, as in measur-

ing the temperature T anisotropies of the CMB, the polarised background needs to be separated from instrumental noise and signals due to astronomical foregrounds (in particular, synchrotron radiation and thermal radiation from Galactic dust). This foreground component separation is more difficult compared with the case of CMB temperature, due to the relative strength and morphological complexity of polarised foregrounds, which are poorly understood.

Foreground component separation has been performed in numerous ways but, on real observational data, always by removing foreground contamination from scalar signals. For example, in the polarised setting, foreground contamination is removed from the Stokes Q and U or from E and B mode maps by treating Q and U or E and B as independent scalar fields. We presented a thorough discussion of blind and non-blind component separation methods in Rogers et al. (2016) (see also e.g., Delabrouille et al. 2009; Bobin et al. 2013 for reviews). In this work, we highlight only the four component separation methods employed in *Planck Collaboration et al. (2015b)*. Commander (Eriksen et al. 2006, 2008) and SEVEM (Martínez-González et al. 2003; Leach et al. 2008; Fernández-Cobos et al. 2012) operate on the Q and U maps, while

* E-mail: keir.rogers.14@ucl.ac.uk

NILC (Delabrouille et al. 2009) and SMICA (Cardoso et al. 2008) operate on the E and B mode maps.

Recently, Fernández-Cobos et al. (2016) explored an extension of the Internal Linear Combination (ILC) method to act fully on the spin-2 signal formed by the Q and U maps and applied their method to simulations. In general, the ILC method estimates the CMB as a weighted sum of maps of the sky at different microwave frequencies. The weights are constrained to conserve the CMB signal but minimise foreground and noise residuals by minimising the variance of the output map. The weights can be localised in various domains, but most usefully in wavelet space (e.g., Rogers et al. 2016), which allows the weights to vary simultaneously with position on the sky and harmonic scale. Fernández-Cobos et al. (2016) minimised the covariant quantity $\langle |P|^2 \rangle$, where $P = Q + iU$, in map space. Consequently, they do not consider any harmonic localisation.

In this work, we introduce the Spin, Scale-discretised, directional wavelet ILC or Spin-SILC. This is an extension of the SILC method we introduced in Rogers et al. (2016), where to analyse spin signals such as CMB polarisation we now use spin scale-discretised wavelets, the complete construction of which is presented in McEwen et al. (2015b) (see also Leistedt et al. 2015; McEwen et al. 2014). Wavelets are functions that are localised in both real and harmonic space and, in particular, scale-discretised wavelets satisfy excellent localisation properties (McEwen et al. 2015a). In the scalar SILC method, we use directional scale-discretised wavelets (Wiaux et al. 2008; McEwen et al. 2013; McEwen et al. 2015b). Directional wavelets are spatially and harmonically localised and additionally “directionally-localised,” i.e., the spatial kernels are non-axisymmetric and can be rotated to pick out a preferred direction on the surface of the sphere. Rogers et al. (2016) gives an introductory summary of directional wavelets. The spin wavelets we use are still spatially, harmonically and directionally localised but are now constructed in the space of spin spherical harmonics. When spin wavelets are convolved with spin signals defined on the sphere, the output wavelet coefficients isolate signal structure of different scale and orientation, while maintaining the spatial information. The spin can in general be arbitrary, but since we are interested in analysing the spin-2 signal $P = Q + iU$ we adopt spin-2 wavelets. By the particular construction of the spin wavelets, the complex spin-2 wavelet coefficients can be separated (by their real and imaginary parts) into scalar wavelet coefficients of the E and B fields, where the scalar wavelet coefficients correspond to a scalar wavelet that is a spin-lowered version of the original spin-2 wavelet (McEwen et al. 2015b). Hence, the spin-2 wavelet transform at the heart of Spin-SILC performs E - B decomposition from input Q and U maps. The ILC method is then applied to the complex wavelet coefficients, with complex weights, and jointly minimises the variance of the reconstructed E and B fields. Moreover, the weights vary spatially, harmonically and according to different orientations, fine-tuning the cleaning algorithm to remove foreground and noise contamination.

It follows that Spin-SILC introduces two main novelties to CMB polarisation component separation. Firstly, the use of spin scale-discretised wavelets allows the full analysis of the polarisation spin signal P . By their construction, we can then perform component separation and E - B decomposition simultaneously and self-consistently. Secondly, the use of directional wavelets allows the additional flexibility to localise the foreground removal according to the morphological structure of the CMB and the foregrounds.

There is a third novel attribute to Spin-SILC of interest to future polarisation observations. Although in this work, we have




Figure 1. The spatial localisation on the sphere of spin, directional, scale-discretised wavelets. The top row shows larger scale wavelets than the bottom row. The left column shows the real part of the wavelet, the middle column shows the imaginary part of the wavelet and the right column shows the absolute value of the wavelet. The number of directions per wavelet scale $N = 5$. Therefore, for complete reconstruction at each scale, the above wavelets would be complemented by four more wavelets of the same size but of a different orientation on the sphere. The spin number $s = 2$, which is what is required for the analysis of Stokes Q and U modes. This figure is adapted from McEwen et al. (2015b).

tested Spin-SILC on the full-sky multifrequency maps provided by the Planck Collaboration, these frequency maps are dominated by instrumental noise and hence so are also our estimates of the CMB polarisation. Future polarisation measurements will have high signal-to-noise, but will usually cover only a fraction of the sky. However, E - B decomposition (from the measured Q and U modes) on the cut-sky is not uniquely defined unlike the full-sky case. This leads to leaking or mixing between the E and B modes. This is of particular concern in extracting the B field since the E field is orders of magnitude larger. As presented in a companion paper (Leistedt et al. 2016), the spin scale-discretised wavelets we use can be employed to construct pure estimators of the masked E and B modes (pure E (B) modes are orthogonal to all B (E) modes on the partial sky, respectively). This builds on the work of Lewis et al. (2002); Bunn et al. (2003); Smith & Zaldarriaga (2007); Grain et al. (2012); Ferté et al. (2013) (see in particular Bunn et al. 2003 for a discussion of pure modes at the map level). This only requires calculating additional wavelet transforms of the input data subject to a suitably apodised mask. One of the main advantages of this approach is the possibility of optimising the mask as a function of scale and direction, therefore yielding a more efficient cancellation of the systematic E - B mixing due to masking (Leistedt et al. 2016). Hence, Spin-SILC can produce accurate estimates of the cosmological E and B fields, even on the cut-sky, in conjunction with the E - B estimators presented in Leistedt et al. (2016). More details of how Spin-SILC can operate on partial sky observations are given in § 3.8.

We provide an introduction to spin scale-discretised wavelets in § 2. In § 3, the Spin-SILC algorithm is explained in detail. We test the method on *Planck* simulations in § 4 and *Planck* data in § 5. In § 6, we compare our method to previous component-separation methods. We discuss the results in § 7 and conclude in § 8.




Figure 2. The harmonic localisation of the spin wavelets used in this work (κ_ℓ^j as defined in Eq. (11)), where j specifies the wavelet scale. Increasing j corresponds to a smaller wavelet kernel and so a multipole range on smaller scales (*i.e.*, larger multipoles ℓ). The largest wavelet scale (Scal.) is the scaling function (§ 3.4). This choice of wavelets deliberately ensures exact reconstruction only for $\ell \leq 2048$. The tapering of the smallest wavelet for $2048 < \ell \leq 2253$ suppresses the smallest-scale power within the algorithm. The band-limits of the above wavelets are given in Table 1.

Table 1. The harmonic band-limits [$\ell_{\min}^j, \ell_{\max}^j$] of the spin wavelets used in this work. ℓ_{peak}^j is the multipole at which each wavelet has its maximum response. The final column shows the number of equiangular samples per wavelet coefficient map N_{samp}^j .

Wavelet scale j	ℓ_{\min}^j	ℓ_{peak}^j	ℓ_{\max}^j	N_{samp}^j
Scal.	0	64	64	8,385
0	32	64	128	33,153
1	64	128	256	131,841
2	128	256	512	525,825
3	256	512	706	998,991
4	542	705	918	1,688,203
5	705	917	1193	2,850,078
6	917	1192	1551	4,815,856
7	1192	1550	2015	8,126,496
8	1550	2015	2253	10,158,778

2 SPIN WAVELETS

Spin, directional, scale-discretised wavelets on the sphere that support exact reconstruction have been constructed in McEwen et al. (2015b) (and discussed briefly in McEwen et al. 2014; Leistedt et al. 2015). These are an extension of the scalar, directional wavelets developed in Wiaux et al. (2008); McEwen et al. (2013), which are used in the scalar version of the SILC method for the analysis of CMB temperature anisotropies (Rogers et al. 2016). They maintain the properties of spatial, harmonic and directional localisation, but can now additionally analyse spin fields by being constructed on the basis of spin spherical harmonics. In particular, spin-2 wavelets can be convolved with the spin-2 field $P = Q + iU$, where Q and U are the Stokes parameters of the CMB’s linear polarisation. Figure 1 shows an example of the spatial localisation of spin wavelets. Unlike scalar wavelets which are real-valued, spin wavelets are complex-valued. Figure 2 shows an example of the harmonic localisation of spin wavelets (for the wavelets used

in this work). Figure 3 shows an example of spin-2 wavelet decomposition as applied to a simulated CMB polarisation field P . By the construction of the spin-2 wavelets, the real and imaginary parts of the complex wavelet coefficient maps of P are respectively scalar wavelet transforms of E and B fields (with a different scalar wavelet). It can be seen that the spin-2 wavelet transform in the Spin-SILC method carries out the decomposition of the CMB polarisation into E and B modes. Details about the use of Spin-SILC on partial sky observations are given in § 3.8.

3 METHOD

We start by outlining the Spin-SILC algorithm. The steps are explained in more detail in the subsequent subsections (§ 3.1 to § 3.7). The use of Spin-SILC on partial sky observations is discussed in § 3.8. We discuss our numerical implementation in § 3.9.

(1) The raw input data are full-sky frequency maps of the anisotropies in the linear polarisation of the CMB, *i.e.*, Stokes Q and U fields. These maps use the HEALPix format (Górski et al. 2005). (See § 3.1.) The model we employ for the raw data is explained in § 3.2.

(2) The maps are “pre-processed” by inpainting in a small point source mask and each convolved to have the same effective beam (see § 3.3).

(3) At each frequency band, the complex spin-2 polarisation field $P = Q + iU$ is formed. Each P map is converted into a set of complex-valued spin-2 wavelet coefficient maps. This separates both the scale and orientation of structure within each map. These wavelet coefficient maps are sampled according to the sampling theorem of McEwen et al. (2015). (See § 3.4.)

(4) A spin-2 ILC method is then applied separately to each wavelet scale and orientation. For each scale and orientation, the multifrequency wavelet coefficient maps are weighted and added to form a single (complex-valued) wavelet coefficient map that contains mainly CMB signal, as well as some residual foreground and noise. These weights are allowed to vary at each wavelet coefficient. The calculation of these weights is explained in § 3.5.

(5) By the construction of the spin-2 wavelets we use, the real and imaginary parts of the final ILC wavelet coefficient maps are respectively ILC estimates of the scalar wavelet transforms of the CMB E and B maps (with a different scalar wavelet). Therefore, the real and imaginary parts are separately synthesised with scalar wavelets to form the final products: full-sky maps of the CMB E and B anisotropies (with some residual foreground and noise). (See § 3.7.) Q and U maps are also formed by a standard spin-2 inverse wavelet transform of the ILC results (see § 3.6). All final maps use the HEALPix format.

3.1 Input data

Our CMB polarisation map products use full-mission 2015 release *Planck* Q and U polarisation maps as their input¹ (Planck Collaboration et al. 2015c,d). All seven polarisation frequency channels are used. At 70 GHz, we use the higher-resolution version at $N_{\text{side}} = 2048$. As noted in Planck Collaboration et al. (2015a), the 100, 143 and 217 GHz polarisation maps have been high-pass filtered due to insufficient characterisation of residual systematic effects on large scales, in particular leakage between temperature and

¹ <http://pla.esac.esa.int/pla>

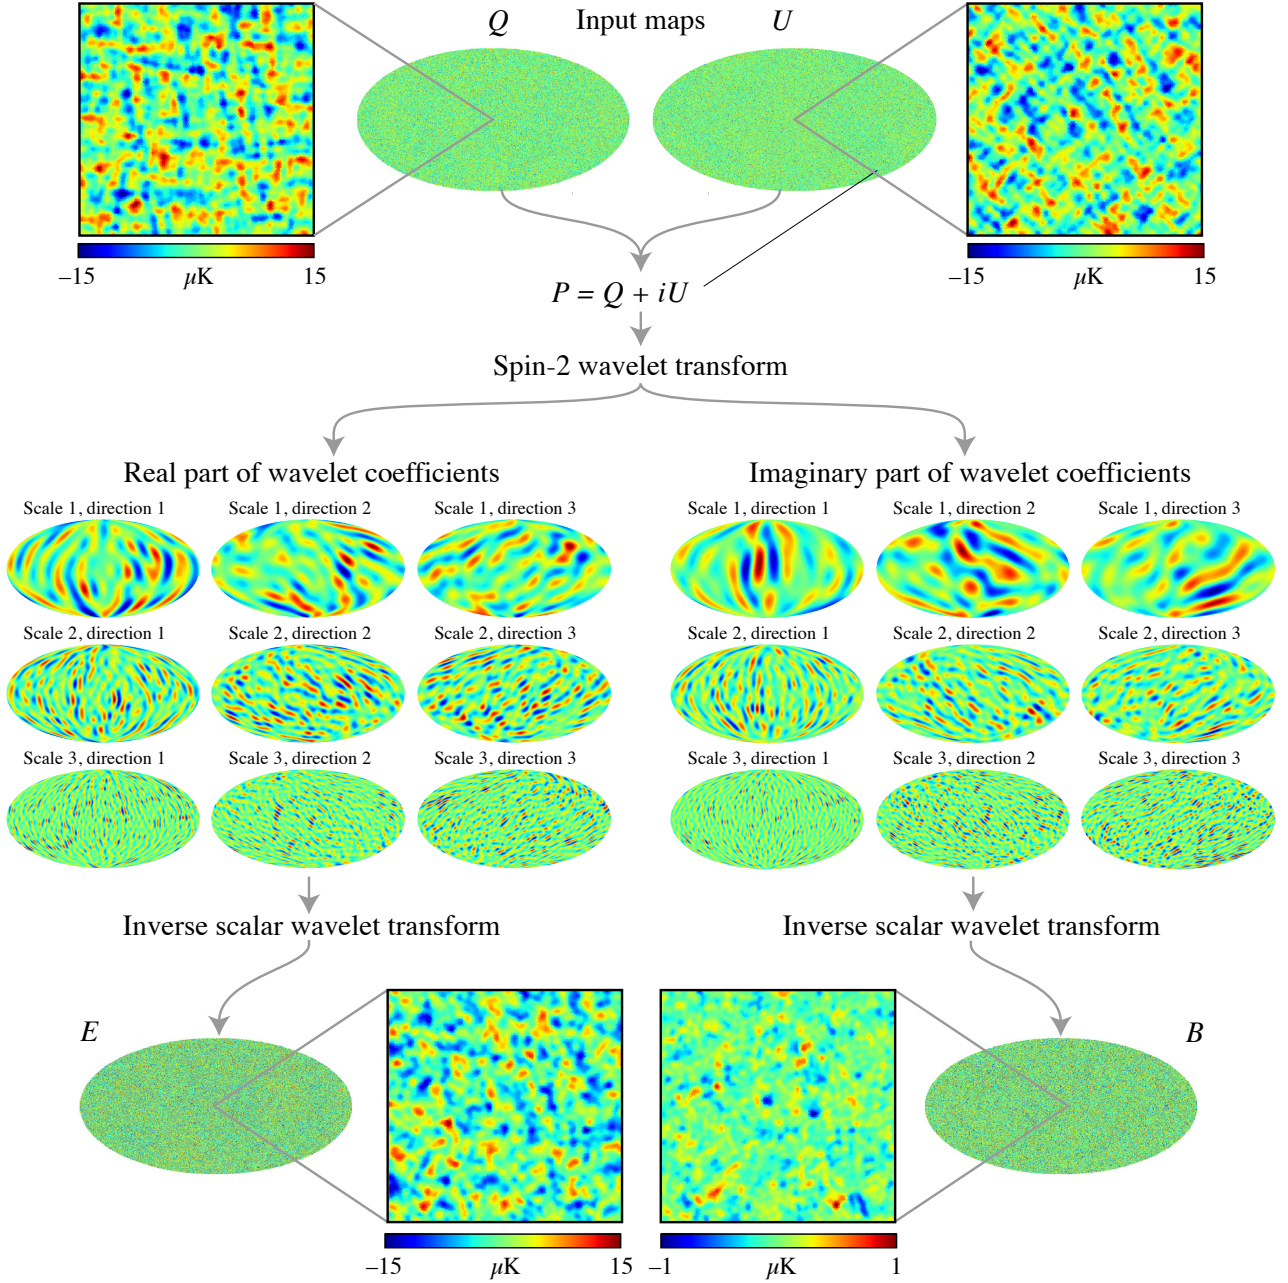


Figure 3. An illustration of the spin, directional wavelet decomposition of the CMB Stokes Q and U maps and the E - B separation that automatically occurs as a consequence. The top row shows example input Q and U maps, simulated with lensed scalar perturbations, with zoomed regions to show structure in the fields. The middle row shows the real and imaginary parts of the spin, directional wavelet coefficient maps, formed by the spin-2 wavelet transform of $P = Q + iU$. The ILC algorithm acts on such wavelet coefficients (calculated for multiple polarisation channels) and produces clean wavelet coefficients of the CMB polarisation. By the construction of the wavelets, the real and imaginary parts are respectively equal to scalar wavelet transforms of the E and B fields (with a different scalar wavelet). The bottom row shows the output E and B maps, also with zoomed regions, formed respectively by inverse scalar wavelet transforms of the real and imaginary parts of the wavelet coefficient maps. In our Spin-SILC analysis we include wavelets on smaller scales than those used in the simple demonstration shown above.

polarisation measurements. We therefore follow [Planck Collaboration et al. \(2015b\)](#) in also high-pass filtering the spherical harmonic coefficients of our output data products with a harmonic cosine fil-

ter:

$$w_\ell = \begin{cases} 0, & \text{if } \ell < 20, \\ \frac{1}{2} \left[1 - \cos\left(\frac{\pi}{20}(\ell - 20)\right) \right], & \text{if } 20 \leq \ell \leq 40, \\ 1, & \text{otherwise.} \end{cases} \quad (1)$$

We use the full-mission Full Focal Plane 8 (FFP8) simulations

(Planck Collaboration et al. 2015e) with lensed scalar perturbations and without bandpass mismatch. These consist of a superposition of a CMB realisation, a noise realisation and full simulations of diffuse and point source astrophysical foregrounds.

While we do not expect an algorithm developed for next generation precision CMB polarisation observations to demonstrate its full capabilities with the *Planck* dataset, this setting comprises the best publicly-available simulations and data, and benefits from the availability of comparison data products from well-studied and highly tested component separation algorithms used by the Planck Collaboration. Thus we use the *Planck* setting to benchmark our algorithm.

3.2 Data model

Each of the full-sky Stokes Q and U polarisation maps ($X = Q, U$) can be independently² physically modelled (e. g., Basak & Delabrouille 2013) as

$$X^{\text{OBS},c}(\hat{n}) = \int_{\hat{n}'} d\hat{n}' B^c(\hat{n}, \hat{n}') X^{\text{SIG},c}(\hat{n}') + X^{\text{N},c}(\hat{n}), \quad (2)$$

where the signal component can further be decomposed as

$$X^{\text{SIG},c}(\hat{n}) = a^c X^{\text{CMB}}(\hat{n}) + X^{\text{FG},c}(\hat{n}). \quad (3)$$

$X^{\text{CMB}}(\hat{n})$ is the CMB component at a point on the sky \hat{n} . $X^{\text{FG},c}(\hat{n})$ and $X^{\text{N},c}(\hat{n})$ are respectively the foreground and detector noise components for frequency channel c . a^c is the calibration coefficient for the CMB for each channel. The overall signal component is smoothed by a beam function $B^c(\hat{n}, \hat{n}')$ due to the finite resolution of the observations. However, the noise component is not smoothed by the beam. Here we assume the beam to be circularly symmetric. Therefore, the beam can be represented as a sum over Legendre polynomials,

$$B^c(\hat{n}, \hat{n}') = \sum_{\ell=0}^{\infty} \frac{2\ell+1}{4\pi} B_{\ell}^c P_{\ell}(\hat{n} \cdot \hat{n}'). \quad (4)$$

We can recast Eq. (2) in the scalar spherical harmonic representation as

$$a_{\ell m}^{\text{OBS},c} = a^c B_{\ell}^c a_{\ell m}^{\text{CMB}} + B_{\ell}^c a_{\ell m}^{\text{FG},c} + a_{\ell m}^{\text{N},c} \quad (5)$$

where $a_{\ell m}$ are the coefficients of scalar spherical harmonics ${}_0Y_{\ell m}(\hat{n})$.

The above is a useful representation for the data pre-processing in § 3.3. However, the novelty of Spin-SILC is to develop a component separation algorithm that directly makes use of the spin properties of the CMB polarisation field. The Stokes Q and U parameters are defined with respect to a fixed coordinate system on the sky. They can be identified by their respective ‘+’ and ‘×’ patterns, as seen in the top row of Fig. 3. However, there is no rotationally-invariant measure of the power as a function of scale. In order to address this, we can first form the complex-valued, spin ± 2 polarisation field

$$\begin{aligned} {}_{\pm 2}P(\hat{n}) &= Q(\hat{n}) \pm iU(\hat{n}) \\ &= \sum_{\ell m} {}_{\pm 2}a_{\ell m} {}_{\pm 2}Y_{\ell m}(\hat{n}). \end{aligned} \quad (6)$$

The spin ± 2 property implies that ${}_{\pm 2}P(\hat{n})$ transforms under local

rotations of angle ψ via ${}_{\pm 2}P \rightarrow e^{-is\psi} {}_{\pm 2}P$, where spin number $s = \pm 2$. The second equality therefore follows by expanding the spin field ${}_{\pm 2}P(\hat{n})$ in the basis of spin spherical harmonics ${}_{\pm 2}Y_{\ell m}(\hat{n})$ with spin spherical harmonic coefficients ${}_{\pm 2}a_{\ell m}$.

We can then define the scalar E and pseudo-scalar B fields in harmonic space as

$$\begin{aligned} {}_0E_{\ell m} &= -\frac{1}{2}({}_2a_{\ell m} + {}_{-2}a_{\ell m}) \\ {}_0B_{\ell m} &= \frac{i}{2}({}_2a_{\ell m} - {}_{-2}a_{\ell m}). \end{aligned} \quad (7)$$

By construction, the E and B fields are real-valued and allow the calculation of rotationally-invariant angular power spectra. E modes are identified by the polarisation strength increasing in a direction parallel or perpendicular to the sense of the polarisation; it is the curl-free component of the spin-2 signal. B modes are identified by the polarisation strength increasing in a direction unaligned to the sense of the polarisation; it is the divergence-free component of the spin-2 signal. E and B modes therefore respectively separate the underlying field into parity-even and parity-odd components.

3.3 Data pre-processing

The input frequency Q and U maps are diffusively inpainted in a small point source mask following the method employed by Planck Collaboration et al. (2015f). This recognises that the ILC fails when the CMB is obscured by bright extragalactic polarised sources. The inpainting removes these sources and replaces them with an extrapolation of the surrounding signal. The mask is the union of the *Planck* LFI and HFI point source masks, which are constructed from the Second Planck Catalogue of Compact Sources (PCCS2) (Planck Collaboration et al. 2015g)³. It masks about 0.6% of the whole sky, predominantly along the Galactic equator.

After the inpainting, we convert all the input frequency maps to the same resolution by performing a deconvolution/convolution procedure that gives scalar spherical harmonic coefficients

$${}_0a_{\ell m}^c = \frac{B_{\ell}^{\text{EFF}}}{B_{\ell}^c} {}_0a_{\ell m}^{\text{OBS},c}, \quad (8)$$

where B_{ℓ}^{EFF} is the beam transfer function giving the resolution at which we perform the ILC. For *Planck* data, we use a Gaussian beam with a FWHM of $5'$ as our input beam. Our final map products are re-convolved to a $10'$ beam in order to suppress residual noise. The beams we deconvolve B_{ℓ}^c are taken from the Reduced Instrument Model (RIMO)⁴. For the LFI beams, we use Gaussian approximations with FWHM $32.33'$, $27.01'$ and $13.25'$ for 30, 44 and 70 GHz respectively. Following Planck Collaboration et al. (2014); Rogers et al. (2016), the deconvolved beams are thresholded such that the B_{ℓ}^c is set to the value given in the RIMO or 0.001, whichever is larger. This suppresses noise within the ILC method.

3.4 Spin wavelet analysis

The spin wavelet ILC method requires the decomposition of each band-limited, complex-valued polarisation map ${}_2P^c(\hat{n})$ (as formed

² The independent modelling is based on the accurate assumption that any mixing of Q and U modes in their measurement has been previously corrected in any given experiment.

³ The details of their construction are given within the FITS files. They can be downloaded from <http://pla.esac.esa.int/pla>.

⁴ *Planck* 2015 Release Explanatory Supplement: The 2015 instrument model (http://wiki.cosmos.esa.int/planckpla2015/index.php/The_RIMO).

by Eq. (6) into a set of spin wavelet coefficient maps $W_p^{2\Psi^j}$. We use the spin, directional, scale-discretised wavelets of [McEwen et al. \(2015b\)](#); [Leistedt et al. \(2015\)](#); [McEwen et al. \(2014\)](#). Following an introductory summary in § 2, we now discuss some of the technical details of our wavelet implementation. We drop the c superscript on ${}_2P(\hat{n})$ for the rest of this subsection since each frequency map is analysed using the same wavelets. We concentrate on the spin-2 wavelet transforms we use in Spin-SILC, but the wavelets we use can be generalised to arbitrary spin.

The spin wavelet coefficients are defined as the directional convolution of ${}_2P$ with spin wavelets ${}_2\Psi^j$ defined on the sphere \mathbb{S}^2 (specifically those shown in Fig. 2), where index j denotes the wavelet scale. Like the scalar case ([Wiaux et al. 2008](#); [McEwen et al. 2013](#)), spin, directional wavelets yield coefficients $W_p^{2\Psi^j}(\hat{p})$ that live on the space of three-dimensional rotations, *i. e.*, the rotation group $\text{SO}(3)$:

$$W_p^{2\Psi^j}(\hat{p}) \equiv \langle {}_2P | \mathcal{R}_{\hat{p}} {}_2\Psi^j \rangle = \int_{\mathbb{S}^2} d\hat{n} {}_2P(\hat{n}) (\mathcal{R}_{\hat{p}} {}_2\Psi^j)^*(\hat{n}), \quad (9)$$

where $d\hat{n}$ is the usual invariant measure on the sphere and $*$ denotes complex conjugation. The rotation operator is defined by

$$(\mathcal{R}_{\hat{p}} {}_2\Psi^j)(\hat{n}) \equiv {}_2\Psi^j(\mathbf{R}_{\hat{p}}^{-1}\hat{n}), \quad (10)$$

where $\mathbf{R}_{\hat{p}}$ is the three-dimensional rotation matrix corresponding to $\mathcal{R}_{\hat{p}}$. In Eqs. (9) and (10), $\hat{p} = (\theta, \phi, \chi) \in \text{SO}(3)$ denotes the Euler angles (in the zyz convention) with colatitude $\theta \in [0, \pi]$, longitude $\phi \in [0, 2\pi]$ and direction $\chi \in [0, 2\pi]$. In other words, the wavelet coefficients probe directional structure in ${}_2P$ with χ corresponding to the orientation about each point (θ, ϕ) on the sphere.

Spin, directional wavelets are defined by their spin spherical harmonic coefficients in factorised form:

$${}_2\Psi_{\ell n}^j \equiv \sqrt{\frac{2\ell+1}{8\pi^2}} \kappa_{\ell 2\zeta_{\ell n}}^j, \quad (11)$$

where κ_{ℓ}^j sets the harmonic localisation (Fig. 2) and ${}_2\zeta_{\ell n}$ sets the directional localisation. Their spin properties are maintained by being built on the basis of spin spherical harmonics. Full details of their construction are given in [McEwen et al. \(2015b\)](#). As in [Rogers et al. \(2016\)](#), we flexibly control the harmonic localisation by using different values of the wavelet dilation parameter λ in different multipole regions. For the wavelets we use in Fig. 2, we use $\lambda = 2, 1.3, 1.1$ with transitions at multipoles $\ell = 512, 2015$. The harmonic bounds of each wavelet for scale j are given by $(\ell_{\min}^j, \ell_{\max}^j) = (\lambda^{j'-1}, \lambda^{j'+1})$, taking account of the different values of λ we use and the stitching-together of wavelets at λ transitions. The index j' refers to the original index of the wavelet scale as if only that single value of λ was used. Their peak response is at $\lambda^{j'}$. The details of our harmonic tiling are given in Table 1. A single parameter N (at all scales) defines the number of directions into which each wavelet scale is localised.

We also use an axisymmetric scaling function ${}_2\Phi$ to form scaling coefficients $W_p^{2\Phi}$ which characterise the largest-scale information (in this work for $\ell < 64$) and live on the sphere. This is motivated by testing in [Rogers et al. \(2016\)](#) that showed that the use of directionality on large scales in the ILC is not effective for CMB reconstruction. These wavelets (and the scaling function) satisfy the standard admissibility criterion for exact reconstruction, *i. e.*, no information is lost in the wavelet and inverse wavelet transforms of a band-limited spin signal. For the chosen band-limit, the smallest wavelet is harmonically-truncated. We choose not to use this wavelet, which means that exact reconstruction is only satisfied

for $\ell \leq 2048$. This allows the tapering of the smallest remaining wavelet (for $2048 < \ell \leq 2253$) to suppress the smallest-scale power in the algorithm.

In order to apply the ILC algorithm, the above continuous wavelet coefficients must be discretised. Since they live on the rotation group $\text{SO}(3)$, we represent them using the sampling scheme of [McEwen et al. \(2015\)](#), which is itself a generalisation of the sampling scheme of [McEwen & Wiaux \(2011\)](#). Since the wavelets are band-limited, we use a multi-resolution scheme where each wavelet scale j is pixellated with a minimal number of samples. This means that each wavelet coefficient map $W_p^{2\Psi^j}$ (band-limited at ℓ_{\max}^j) is only evaluated at samples $(\theta_t^j, \phi_p^j, \chi_n)$, where $t \in \{0, 1, \dots, \ell_{\max}^j\}$, $p \in \{0, 1, \dots, 2\ell_{\max}^j\}$ and $n \in \{0, 1, \dots, N-1\}$. In this way, each spin wavelet coefficient map can be separated into N spin, directional wavelet coefficient maps W_{jnk}^{2P} according to the value of n , where k indexes pixel number according to samples (θ_t^j, ϕ_p^j) on the sphere. It follows that each input frequency P map has been decomposed into wavelet coefficient maps, each localised according to harmonic scale j and orientation of structure n , while maintaining spatial localisation (pixel number k). (See Fig. 3 for a demonstration of this decomposition.)

3.5 ILC method

Following the spin, directional wavelet analysis of the input P maps (see § 3.4), there is a spin, directional wavelet coefficient map $W_{jnk}^{2P,c}$ for each channel c , scale j and orientation n with a pixel index k . With this compact notation, we conflate the scaling coefficient maps with the wavelet coefficient maps as the ILC method applies in exactly the same way. We develop the spin wavelet ILC method by an extension of the scalar wavelet ILC method we developed in [Rogers et al. \(2016\)](#) to operate on the complex-valued wavelet coefficient maps we now have. Similar to [Fernández-Cobos et al. \(2016\)](#), we consider the complex spin signal ${}_2P$ rather than considering scalar fields independently (*e. g.*, Q and U independently or E and B independently). However, unlike [Fernández-Cobos et al. \(2016\)](#), who work jointly on Q and U maps in real space, we work in wavelet space, where spatial, scale and directional localisation is possible. The most general extension of the scalar ILC is to estimate the CMB at each wavelet scale and orientation as a sum of wavelet coefficient maps for each frequency with complex-valued weights ω_{jnk}^c :

$$W_{jnk}^{2P,\text{ILC}} \equiv \sum_{c=1}^{N_c} \omega_{jnk}^c W_{jnk}^{2P,c}, \quad (12)$$

where N_c is the number of input channels.

In order to recover an unbiased estimate of the CMB, we impose a constraint on the weights such that

$$\sum_{c=1}^{N_c} a^c \omega_{jnk}^c = 1 + 0i, \quad (13)$$

where we remind the reader that a^c is the real-valued set of calibration coefficients for the CMB Q and U maps introduced in Eq. (3). In order to calculate the weights at each pixel k , we choose to minimise the covariant quantity $\left| W_{jnk}^{2P,\text{ILC}} \right|^2$ with respect to the complex-valued weights ω_{jnk}^c themselves, under the constraint in Eq. (13). This minimisation can be carried out with complex Lagrange multipliers (similarly to the scalar case) giving complex-

valued weights

$$\omega_{jnk}^c = \frac{\sum_{c'=1}^{N_c} (R_{jnk}^{-1})^{cc'} a^{c'}}{\sum_{c=1}^{N_c} \sum_{c'=1}^{N_c} a^c (R_{jnk}^{-1})^{cc'} a^{c'}}, \quad (14)$$

where the true covariance matrices at scale j , orientation n and pixel k , $(R_{jnk})^{cc'} = \langle W_{jnk}^{2P,c*} W_{jnk}^{2P,c'} \rangle$ (where the angled brackets indicate an ensemble average, although in practice we empirically estimate these covariances as explained below). In this work, we assume $a^c = 1, \forall c$, i. e., that the CMB is perfectly calibrated in the data we use.

There are two main consequences from minimising the quantity we choose. First, as in scalar SILC, we assume that the CMB and foregrounds and the CMB and noise are respectively uncorrelated. It follows that the ensemble cross-term between CMB and residual contamination is zero and, since the CMB is conserved by the constraint in Eq. (13), we are minimising only the variance of the error in CMB reconstruction. Second, by minimising the variance of the full complex-valued spin-2 wavelet coefficients using complex weights, we are in turn jointly minimising the variance of the E and B modes. This is thanks to the construction of the spin wavelets we use, as discussed in § 3.6. Unlike a foreground cleaning algorithm acting on the Q and U or E and B maps separately, this approach ensures that all the information (i. e., from the multiple polarisation channels and the Q and U cross terms) is used to jointly construct clean CMB polarisation P and its E and B modes.

As in scalar SILC, we estimate the covariance matrices $(R_{jnk})^{cc'}$ empirically on the data. We achieve this by replacing the appropriate ensemble average with a weighted average of the surrounding pixels. Specifically, we smooth the maps of covariance matrix elements with a Gaussian kernel in harmonic space. The size of this kernel is proportional to the size of the wavelet used at each scale⁵. Full details of this empirical estimation of covariances and possible optimisations to the method are given in Rogers et al. (2016).

3.6 Spin wavelet synthesis to Stokes Q and U modes

Although a novelty of Spin-SILC is to simultaneously perform E - B decomposition and component separation (i. e., to synthesise the ILC results directly to E and B maps as explained in § 3.7), one can also form Q and U maps. This is carried out by a single spin-2 inverse wavelet transform. After the ILC method (see § 3.5) has been applied to the frequency wavelet coefficient maps, there is one ILC estimate of the CMB P field (with some residual foreground and noise) at each wavelet scale and orientation $W_{jnk}^{2P,ILC}$. Multiple orientations $\chi_0, \chi_1, \dots, \chi_{N-1}$ are combined at each scale to form wavelet coefficient maps $W_P^{2\Psi^j,ILC}(\hat{\rho})$ that live on $SO(3)$. We also have the scaling coefficient map $W_P^{2\Phi,ILC}(\hat{n})$ that lives on the sphere (and characterises the largest scales). In order to calculate our real space estimate of the CMB polarisation spin field ${}_2P^{ILC}(\hat{n})$ (and hence the Stokes parameters $Q^{ILC}(\hat{n})$ and $U^{ILC}(\hat{n})$),

we perform the following spin-2 inverse wavelet transform:

$$\begin{aligned} {}_2P^{ILC}(\hat{n}) &= Q^{ILC}(\hat{n}) + iU^{ILC}(\hat{n}) \\ &= \int_{\mathbb{S}^2} d\hat{n}' W_P^{2\Phi,ILC}(\hat{n}') (\mathcal{R}_{\hat{n}' \hat{n}} \Phi)(\hat{n}) \\ &\quad + \sum_j \int_{SO(3)} d\hat{\rho} W_P^{2\Psi^j,ILC}(\hat{\rho}) (\mathcal{R}_{\hat{\rho} \hat{n}} \Psi^j)(\hat{n}), \end{aligned} \quad (15)$$

where $d\hat{\rho}$ is the usual invariant measure on the rotation group. We have used the same spin wavelets as in the wavelet analysis in § 3.4. The final ILC Q and U maps are pixellated in the HEALPix format.

3.7 Scalar wavelet synthesis to E and B modes

A considerable advantage of the Spin-SILC method is that it simultaneously removes foreground and noise contamination from the cosmological signal and carries out the E - B decomposition discussed in § 3.2. It achieves the latter by using a property of the spin scale-discretised wavelets that relates $W_P^{2\Psi^j}(\hat{\rho})$, the spin-2 wavelet transform of P to $W_E^{0\Psi^j}(\hat{\rho})$ and $W_B^{0\Psi^j}(\hat{\rho})$, the scalar wavelet transforms of \tilde{E} and \tilde{B} :

$$\begin{aligned} W_E^{0\Psi^j}(\hat{\rho}) &= -\text{Re}[W_P^{2\Psi^j}(\hat{\rho})] \\ W_B^{0\Psi^j}(\hat{\rho}) &= -\text{Im}[W_P^{2\Psi^j}(\hat{\rho})]. \end{aligned} \quad (16)$$

The intermediate fields \tilde{E} and \tilde{B} are respectively related to E and B by a harmonic normalisation of their scalar spherical harmonic coefficients:

$$\begin{aligned} {}_0E_{\ell m} &= \frac{1}{N_{\ell,2}} {}_0\tilde{E}_{\ell m} \\ {}_0B_{\ell m} &= \frac{1}{N_{\ell,2}} {}_0\tilde{B}_{\ell m}, \end{aligned} \quad (17)$$

where $N_{\ell,s} = \sqrt{\frac{(\ell+s)!}{(\ell-s)!}}$.

The straightforward E - B decomposition is achieved by the construction of the wavelets and is discussed in detail in McEwen et al. (2015b) and Leistedt et al. (2016). In Eq. (16) the scalar wavelets ${}_0\Psi^j$ are spin-lowered versions of the spin-2 wavelets ${}_2\Psi^j$:

$${}_0\Psi^j(\hat{n}) = \tilde{\delta}^{-2} {}_2\Psi^j(\hat{n}). \quad (18)$$

(An equivalent equation links the scalar and spin scaling functions.) $\tilde{\delta}$ is a first-order differential operator known as the spin-lowering operator since it lowers the spin of spherical harmonic functions: $\tilde{\delta}_s Y_{\ell m}(\hat{n}) = \frac{1}{N_{\ell,s}} {}_{s-1}Y_{\ell m}(\hat{n})$.

By applying Eq. (16), we can separate the spin wavelet coefficient maps $W_P^{2\Psi^j,ILC}(\hat{\rho})$ (defined in § 3.6) into scalar wavelet coefficient maps of the intermediate \tilde{E} and \tilde{B} modes $W_Y^{0\Psi^j,ILC}(\hat{\rho})$, where $Y = \tilde{E}, \tilde{B}$. An equivalent separation forms the scaling coefficient maps $W_Y^{2\Phi,ILC}(\hat{n})$. In order to calculate our real space estimates of the CMB \tilde{E} and \tilde{B} modes ${}_0Y^{ILC}(\hat{n})$, we perform inverse scalar wavelet transforms with the (spin-lowered) scalar scaling function ${}_0\Phi$ and scalar wavelets ${}_0\Psi^j$ (as defined in Eq. (18)):

$$\begin{aligned} {}_0Y^{ILC}(\hat{n}) &= \int_{\mathbb{S}^2} d\hat{n}' W_Y^{2\Phi,ILC}(\hat{n}') (\mathcal{R}_{\hat{n}' \hat{n}} \Phi)(\hat{n}) \\ &\quad + \sum_j \int_{SO(3)} d\hat{\rho} W_Y^{0\Psi^j,ILC}(\hat{\rho}) (\mathcal{R}_{\hat{\rho} \hat{n}} \Psi^j)(\hat{n}). \end{aligned} \quad (19)$$

The output scalar spherical harmonic coefficients can be renormalised to the usual E and B fields by applying Eq. (17). The final ILC E and B maps are pixellated in the HEALPix format.

⁵ $\text{FWHM}^j = 50 \sqrt{\frac{1200}{N_{\text{sam}}^j}}$. This value is the same as used in the NILC implementation on *Planck* data.




Figure 4. Planck simulations. Differences between output ILC and input CMB maps from FFP8 simulations with lensed scalar perturbations. The maps have been smoothed to $\text{FWHM} = 80'$ and downgraded to $N_{\text{side}} = 128$. The grey pixels are the UPB77 confidence mask from [Planck Collaboration et al. \(2015b\)](#), which masks the Galactic region in FFP8 simulations where foreground emission is strongest. From top to bottom, we show differences in (a) Stokes Q and (b) Stokes U maps.

3.8 Spin-SILC on partial sky observations

We have outlined above the Spin-SILC method specifically as it applies on the full sky. However, Spin-SILC is being primarily developed for application to future CMB polarisation experiments, which will have greater signal-to-noise and/or resolution, but will typically observe only part of the sky. The decomposition of Stokes Q and U measurements into E and B modes is essential for cosmological analyses, in particular for a measurement of the BB angular power spectrum. This is strictly well-defined only on the whole sky, as in Eq. (7).

This decomposition is not well-defined if the input measurements only cover a part of the sky. However, following [Bunn et al. \(2003\)](#), a polarisation field on the cut-sky can be decomposed into a complete orthonormal basis defined by “pure E ”, “pure B ” and “ambiguous” modes. Pure E modes have vanishing curl and are orthogonal to all B modes on the partial sky. Pure B modes have vanishing divergence and are orthogonal to all E modes on the partial sky. Ambiguous modes are all other modes, which will have both vanishing divergence and curl. It is the inability to distinguish ambiguous modes which leads to the problem of E - B leakage where ambiguous modes are erroneously counted as E or B . However, following, *e.g.*, [Bunn et al. \(2003\)](#); [Smith & Zaldarriaga \(2007\)](#), if the pure B modes can be isolated, they will form an estimate of the cosmological B power, unbiased by E - B leaking.




Figure 5. Planck simulations. From top to bottom, (a) EE , (b) BB and (c) EB angular power spectra comparing output ILC in the axisymmetric limit ($N = 1$) to input CMB from FFP8 simulations with lensed scalar perturbations. In the top panel (a), the thin red line shows residuals after subtracting the input CMB spectrum.

[Leistedt et al. \(2016\)](#) have shown how the spin wavelets we use can be employed to construct estimates of the pure modes defined above on a masked sky. This builds on the work of, *e.g.*, [Lewis et al. \(2002\)](#); [Smith & Zaldarriaga \(2007\)](#); [Grain et al. \(2012\)](#). In particular, the spin wavelet pure-mode estimation of [Leistedt et al. \(2016\)](#) requires only two additional wavelet transforms of the input Q and U data and a suitably apodised mask. The Spin-SILC method can be applied on partial sky observations by coherently combining the full-sky method with the pure mode estimation. The application of Spin-SILC to partial sky observations will be investigated in future work, providing the first integrated pipeline to simultaneously carry out E - B decomposition and foreground component separation for future CMB polarisation experiments.

3.9 Numerical implementation

Spin-SILC is implemented in Python and is parallelised. At full Planck resolution ($N_{\text{side}} = 2048$, $\ell_{\text{max}} = 2253$), when run on a 60-core symmetric multiprocessor (SMP) with 1.5 TB RAM and




Figure 6. *Planck* simulations. Differences between angular power spectra of different values of N minus the axisymmetric limit ($N = 1$). The input data are FFP8 simulations with lensed scalar perturbations. *From top to bottom*, we show differences in (a) EE and (b) BB spectra. We note the small amplitude of the reductions in reconstruction residuals from increasing N .

a 24-core cluster node with 256 GB RAM⁶, the pipeline takes approximately 1.5 hours per direction. The wavelet transforms in Spin-SILC are carried out using the latest version of the S2LET⁷ code (Leistedt et al. 2013; McEwen et al. 2015), written in C with Python wrappers. This employs SSHT⁸ (McEwen & Wiaux 2011) and SO3⁹ (McEwen et al. 2015) to compute spin spherical harmonics and Wigner transforms exactly and efficiently. Spin-SILC is developed from the scalar SILC¹⁰ code (Rogers et al. 2016) (which performs component separation on the temperature anisotropies of the CMB).

4 APPLICATION TO PLANCK SIMULATIONS

We tested Spin-SILC on the fiducial full-mission *Planck* FFP8 simulated Stokes Q and U sky maps. We use simulations with lensed scalar perturbations. Figure 4 shows the differences between the reconstructed CMB (using $N = 1$) and the input simulated CMB. The two panels show the differences in Q and U maps, as this most directly compares to the input data. The most striking features are the reductions in residuals in the top left and bottom right corners, aligning with the Ecliptic poles. This reflects reduced noise residuals because there is less noise in the input data due to the scanning strategy of the *Planck* satellite, which integrated for longer

⁶ The exact specification for our infrastructure is an Intel Xeon E7-4890 2.8 GHz SMP with 4×15 -core CPUs with 25.6 GB RAM per core, and an Intel Xeon E5-2697 2.7 GHz node with 2×12 -core CPUs with 10.7 GB RAM per core.

⁷ <http://www.s2let.org>

⁸ <http://www.spinsht.org>

⁹ <http://www.sothree.org>

¹⁰ <http://www.silc-cmb.org>




Figure 7. *Planck* data. *From left to right*, (a) the CMB polarisation E map and (b) the CMB polarisation B map reconstructed using Spin-SILC in the axisymmetric limit ($N = 1$, FWHM = $10'$, $N_{\text{side}} = 1024$).

in those directions. In general, the difference maps are consistent with noise residuals. This is a consistent attribute of the *Planck* polarisation datasets. Figure 5 compares the full-sky angular power spectra ($D_\ell = \ell(\ell + 1)C_\ell/2\pi$) of the same reconstructed CMB and the input signal. The three panels respectively compare the EE , BB and EB spectra, as these are the cosmologically-interesting observables. The EE and BB spectra are consistent with significant residual noise power due to the noisiness of the input maps, although the first four acoustic peaks of the EE spectrum are discernible nonetheless. The reconstructed EB spectrum is consistent with the zero input value.

We tested the impact of using directional spin wavelets in the Spin-SILC method with the simulated dataset. Figure 6 compares the differences in full-sky power spectra between the directional case (for $N = 5, 10$) minus the axisymmetric limit ($N = 1$). The two panels compare EE and BB spectra. It can be seen that using more directional wavelets per wavelet scale reduces power spectrum reconstruction residuals with respect to the axisymmetric limit, very modestly on large scales and more so on small scales. However, the magnitude of these reductions is very small compared to the total power in the output ILC maps, which are dominated by residual instrumental noise; in the BB spectrum, the reduction is comparable to the magnitude of the input lensing signal. These results are fully expected following Rogers et al. (2016), where it was found that the gains in component separation efficacy from employing directionality was marginal in the low signal-to-noise (S/N) regime.

5 APPLICATION TO PLANCK DATA

After testing Spin-SILC on the simulated dataset, we apply as input data the real full-mission *Planck* Stokes Q and U maps. Figure 7 shows our main output data products: full-sky ILC estimates of the CMB polarisation E and B modes using Spin-SILC in the axisymmetric limit ($N = 1$). We show E and B maps in order to highlight the E - B decomposition from input Q and U maps that Spin-SILC automatically carries out thanks to the construction of the spin wavelets that we use (see § 3.4). We reiterate that these maps have been high-pass filtered (for $\ell < 40$) in order to mitigate for residual systematics in the *Planck* polarisation data (see § 3.1). The maps are consistent with large levels of residual instrumental noise, with the scanning pattern of the *Planck* satellite clearly visible. We also note the poor reconstruction in the Galactic plane, particularly towards the Galactic centre, where foreground emission is strongest and most complex.




Figure 8. *Planck* data. Differences between the axisymmetric limit ($N = 1$) of Spin-SILC, NILC and SMICA. The maps have been smoothed to $\text{FWHM} = 80'$ and downgraded to $N_{\text{side}} = 128$. The grey pixels are the UPB77 confidence mask from [Planck Collaboration et al. \(2015b\)](#), which masks the regions of the NILC and SMICA maps not recommended for cosmological analysis. The differences are (from top to bottom) (a) SILC ($N = 1$) - NILC, (b) SILC ($N = 1$) - SMICA and (c) NILC - SMICA; and in (from left to right) (i) Stokes Q and (ii) Stokes U maps.

6 COMPARISON TO PREVIOUS WORK

Having presented the main results of applying Spin-SILC to *Planck* data, we can perform a validation check by comparing to other component separation reconstructions of the CMB from the same dataset. For this purpose, we concentrate on the methods NILC ([Delabrouille et al. 2009](#)) and SMICA ([Cardoso et al. 2008](#)), which are two of the four methods used internally by the Planck Collaboration ([Planck Collaboration et al. 2015b](#)): the former because it is the most similar method to Spin-SILC and the latter because it is the baseline method adopted by the Planck Collaboration for high-resolution analyses. Like Spin-SILC, NILC is an internal linear combination (ILC) method performed in wavelet space. Unlike Spin-SILC, NILC uses scalar axisymmetric wavelets, specifi-

cally scalar needlets ([Narcowich et al. 2006](#); [Marinucci et al. 2008](#); [Baldi et al. 2009](#)), rather than the spin directional wavelets we use ([McEwen et al. 2015b](#); [Leistedt et al. 2015](#); [McEwen et al. 2014](#)) (although spin needlets ([Geller et al. 2008](#)) and mixed needlets ([Geller & Marinucci 2011](#)) have also been developed). This means that in its extension to polarisation ([Basak & Delabrouille 2013](#); [Planck Collaboration et al. 2015b](#)), NILC acts independently on input E and B maps, having been previously decomposed from the original Stokes Q and U data. Similarly to ILC methods, SMICA forms a linear combination of multifrequency data, but in harmonic space. Unlike blind ILC methods which require no physical modelling of the sky components, SMICA is only semi-blind in that on large scales, rather than empirically estimating covariances on the




Figure 9. Planck data. From top to bottom, (a) EE , (b) BB and (c) EB angular power spectra comparing the axisymmetric limit ($N = 1$) of Spin-SILC to NILC and SMICA. In the top panel (a), the thin lines show residuals after subtracting the best-fit Λ CDM model from the *Planck* 2015 likelihood.

data (as on small scales), a fit is performed to a model of the component covariances, with the option to constrain these covariances. This is extended to polarisation by performing a joint processing of the E and B modes in harmonic space. A further difference between ILC methods and SMICA is that SMICA has no spatial localisation in its component separation, although a wavelet implementation of SMICA does exist (Moudden et al. 2005). Indeed, Spin-SILC localises with regard to the greatest number of domains of information, allowing spatial, harmonic and morphological localisation through the use of directional wavelets (see Rogers et al. 2016 for a discussion of the morphological localisation properties of directional wavelets in the SILC method). A significant advantage of Spin-SILC over existing component separation methods (including NILC and SMICA) is the use of spin wavelets, which allows simultaneous component separation and E - B decomposition.

We can empirically compare the three methods with an analysis of the CMB maps reconstructed from the (full-mission 2015 release) *Planck* data and full-sky power spectra measured from those maps. Figure 8 shows the differences between the CMB reconstructed by Spin-SILC (in the axisymmetric limit $N = 1$), NILC and SMICA. We show differences in Q and U maps as this most directly compares to the map products provided by the Planck Collaboration. The differences between the three methods are small in

magnitude in both Q and U and mostly concentrated at the edges of the Galactic mask towards the Galactic centre, where foreground emission is most intense and complex. Quantitatively, we can compare the mean values and standard deviations of the full-sky difference maps. The mean values of the Q difference maps in Figs. 8 (i) (a), (b) and (c) (from top to bottom on the left-hand side) are respectively 5.2×10^{-5} , 4.3×10^{-5} and -9.3×10^{-6} μ K, while the standard deviations are 0.34, 0.37 and 0.34 μ K². The mean values of the U difference maps in Figs. 8 (ii) (a), (b) and (c) (from top to bottom on the right-hand side) are respectively 6.4×10^{-5} , -6.9×10^{-5} and -1.3×10^{-4} μ K, while the standard deviations are 0.33, 0.37 and 0.31 μ K². These values are small and similar, suggesting a strong consistency between the three methods. As discussed in § 3.1, some of the input *Planck* data are high-pass filtered and so are also the output results of all three methods (for $\ell < 40$). This means that any comparison can only be carried out for $\ell \geq 40$.

Figure 9 compares full-sky power spectra measured from component separation maps with CMB spectra derived from the *Planck* 2015 TT and low TEB likelihood¹¹. The three panels respectively compare EE , BB and EB spectra. As with the simulated results in § 4, the EE and BB spectra from all three methods are consistent with significant residual noise power due to the noisiness of the input maps. The only discernible difference is marginally less power in the SMICA maps at multipoles around $\ell = 250$. This could be attributed to the semi-blindness of SMICA better characterising the noise properties of the *Planck* data. The two blind methods, Spin-SILC and NILC have near-identical spectra at all multipoles. The EB spectra of all three methods are consistent with zero.

The comparison of Spin-SILC to existing methods NILC and SMICA has strongly validated the results we showed in § 5. An analysis of maps and power spectra shows an internal consistency between the three algorithms. It also shows that the power of Spin-SILC and other component separation methods is limited by the low S/N of the *Planck* polarisation data, with large amounts of residual noise in the reconstructed CMB. The full potential of the Spin-SILC method thus awaits the input of higher-S/N polarisation data available from upcoming CMB observations.

7 DISCUSSION

The testing of Spin-SILC on *Planck* simulations in § 4 and data in § 5 shows that the use of spin wavelets in CMB polarisation component separation can successfully reconstruct the cosmological background. This is particularly true of the EE power spectrum with the clear detection of the first four acoustic peaks in both simulations and real data. The residual maps to the input simulated CMB (Fig. 4) and power spectra estimated from the SILC maps (Fig. 5) show high levels of residual noise, reflecting the relatively low S/N of the *Planck* data we used. In § 6, we carried out a comparison of the Spin-SILC method with two of the most accurate existing component separation algorithms, NILC and SMICA. We validated our main results by showing a strong internal consistency in reconstructed CMB maps (Fig. 8) and power spectra measured from those maps (Fig. 9). However, this comparison also revealed high

¹¹ The parameters come from the `base_plikHM.TT.lowTEB` likelihood. The values are available in the *Planck* 2015 Release Explanatory Supplement: 2015 Cosmological parameters and MC chains (http://wiki.cosmos.esa.int/planckpla2015/images/f/f7/Baseline_params_table_2015_limit68.pdf).

levels of residual noise in the CMB estimated by all three methods, due to the S/N limitations of the *Planck* polarisation data.

We also tested the use of directional spin wavelets in the Spin-SILC method on the simulations in § 4. We found very modest reductions in reconstructed residual power (Fig. 6) as the amount of directionality was increased. However, given the low S/N input data, the magnitude of these reductions is much smaller than the overall amount of residual power, though at the accuracy required to reconstruct the cosmological *BB* signal, this level of power reduction may become relevant in the high S/N regime. As discussed in Rogers et al. (2016) in the scalar SILC method, instrumental noise has no particular directional structure and thus in the low S/N regime the use of directionality is expected to only have a small effect on the estimate of the reconstructed CMB. The community is only just beginning to accurately characterise polarised foregrounds at high resolution at a range of frequencies. If the foregrounds are complex in high S/N observations, the ability to use directional wavelets may prove useful in localising component separation according to the morphology of the CMB and foregrounds.

Spin-SILC introduces a number of novelties into CMB polarisation component separation. Most notably, the use of spin wavelets allows simultaneous *E-B* decomposition and the joint minimisation of *E* and *B* auto-correlations in residual contamination. Moreover, the use of directionality allows the fine-tuning of the cleaning algorithm according to the morphology of the local signal. As discussed in Rogers et al. (2016), there are various sources of error in the ILC method, which will affect the spin-ILC in an equivalent fashion. Of particular note is the ILC bias, corresponding to the empirical cancellation of CMB modes due to chance correlations with foregrounds and noise (see Delabrouille et al. 2009 for a fuller discussion of this effect), which will also affect the reconstruction of the CMB polarisation. The amount of cancellation may increase with the amount of directionality used within the method. In Rogers et al. (2016), we showed that this can be mitigated either directly from the data or through suites of simulations.

8 CONCLUSIONS

We have presented Spin-SILC, a foreground component separation method specifically developed for the analysis of CMB polarisation data. The use of spin wavelets allows the full analysis of the spin-2 polarisation signal $P = Q + iU$, formed by the Stokes *Q* and *U* parameters. By the particular construction of the spin wavelets we use, Spin-SILC carries out the decomposition of the polarisation signal into *E* and *B* modes by separating the real and imaginary parts of the complex spin-2 wavelet coefficients. This occurs simultaneously to the component separation, where the auto-correlations of *E* and *B* modes are jointly minimised in residual contamination to the reconstructed CMB. Moreover, the wavelets we use are directional. This allows different directional morphologies of CMB and polarised foreground to be separated. This extra information can then be used to better localise the Spin-SILC cleaning algorithm.

We have tested Spin-SILC on full-mission *Planck* simulations and data. We showed that the method can accurately extract cosmological information from input *Q* and *U* maps. We also validated our main results with a comparison to the internal *Planck* methods, NILC and SMICA, showing a strong consistency in both CMB maps and power spectra, with small residuals compared to the two. However, we note that the analysis in this paper is limited by the low S/N of the *Planck* polarisation data. Our final *E* and *B* maps (as well as those of NILC and SMICA) are domi-

nated by residual instrumental noise. Moreover, the full power of the use of directionality in Spin-SILC cannot be fully explored due to the high level of noise in the *Planck* input data. If polarised foregrounds have complex morphology in the high S/N regime, then the use of directionality may prove a useful extra tool in extracting the CMB. In general, it will be interesting to test Spin-SILC further with the high S/N data of upcoming CMB polarisation observations. We make our *Q*, *U*, *E* and *B* maps available at <http://www.silc-cmb.org>¹².

Furthermore, Spin-SILC can be combined with the estimators of Leistedt et al. (2016) to perform component separation on the cut-sky and give accurate estimates of pure *E* and *B* modes (pure *E* (*B*) modes are orthogonal to all *B* (*E*) modes on the cut-sky, respectively). It achieves this in a straightforward fashion (with only two additional wavelet transforms of the input data) due to the construction of the spin wavelets (see Leistedt et al. 2016 for more details about pure mode estimation on the cut-sky using spin wavelets). This is of particular importance for the upcoming high resolution, high S/N CMB polarisation experiments, which will typically make partial-sky observations. Spin-SILC will provide a computationally-efficient algorithm to perform simultaneous *E-B* decomposition and accurate foreground component separation for these next-generation experiments.

ACKNOWLEDGEMENTS

KKR thanks Franz Elsner and Stephen Feeney for valuable discussions. KKR was supported by the Science and Technology Facilities Council. HVP and BL were partially supported by the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 306478-CosmicDawn. JDM was partially supported by the Engineering and Physical Sciences Research Council (grant number EP/M011852/1). AP was supported by the Royal Society. Based on observations obtained with *Planck* (<http://www.esa.int/planck>), an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA and Canada.

REFERENCES

- Baldi P., Kerkycharian G., Marinucci D., Picard D., 2009, *Ann. Stat.*, 37 No.3, 1150, [arXiv:math/0606599](https://arxiv.org/abs/math/0606599)
- Basak S., Delabrouille J., 2013, *MNRAS*, 435, 18, [1204.0292](https://arxiv.org/abs/1204.0292)
- Bobin J., Starck J.L., Sureau F., Basak S., 2013, *A&A*, 550, A73, [1206.1773](https://arxiv.org/abs/1206.1773)
- Bunn E.F., Zaldarriaga M., Tegmark M., de Oliveira-Costa A., 2003, *Phys. Rev. D*, 67, 2, 023501, [astro-ph/0207338](https://arxiv.org/abs/astro-ph/0207338)
- Cardoso J.F., Martin M., Delabrouille J., Betoule M., Patanchon G., 2008, *ArXiv e-prints*, [0803.1814](https://arxiv.org/abs/0803.1814)
- Delabrouille J., Cardoso J.F., Le Jeune M., Betoule M., Fay G., Guillaux F., 2009, *A&A*, 493, 835, [0807.0773](https://arxiv.org/abs/0807.0773)
- Eriksen H.K., Jewell J.B., Dickinson C., Banday A.J., Górski K.M., Lawrence C.R., 2008, *ApJ*, 676, 10, [0709.1058](https://arxiv.org/abs/0709.1058)
- Eriksen H.K., et al., 2006, *ApJ*, 641, 665, [astro-ph/0508268](https://arxiv.org/abs/astro-ph/0508268)
- Errard J., Feeney S.M., Peiris H.V., Jaffe A.H., 2015, *ArXiv e-prints*, [1509.06770](https://arxiv.org/abs/1509.06770)

¹² The DOI for our data release is 10.5281/zenodo.50579.

- Fernández-Cobos R., Marcos-Caballero A., Vielva P., Martínez-González E., Barreiro R.B., 2016, ArXiv e-prints, [1601.01515](#)
- Fernández-Cobos R., Vielva P., Barreiro R.B., Martínez-González E., 2012, MNRAS, 420, 2162, [1106.2016](#)
- Ferté A., Grain J., Tristram M., Stompor R., 2013, Phys. Rev. D, 88, 2, 023524, [1305.7441](#)
- Geller D., Hansen F.K., Marinucci D., Kerkycharian G., Picard D., 2008, PRD, 78, 12, 123533, [arXiv:0811.2881](#)
- Geller D., Marinucci D., 2011, J. Math. Anal. Appl., 375, 2, 610, [arXiv:1006.3835](#)
- Górski K.M., Hivon E., Banday A.J., Wandelt B.D., Hansen F.K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759, [astro-ph/0409513](#)
- Grain J., Tristram M., Stompor R., 2012, Phys. Rev. D, 86, 7, 076005, [1207.5344](#)
- Kamionkowski M., Kosowsky A., Stebbins A., 1997, Physical Review Letters, 78, 2058, [astro-ph/9609132](#)
- Leach S.M., et al., 2008, A&A, 491, 597, [0805.0269](#)
- Leistedt B., McEwen J.D., Büttner M., Peiris H.V., 2016, ArXiv e-prints, [1605.01414](#)
- Leistedt B., McEwen J.D., Büttner M., Peiris H.V., Vanderghelynst P., Wiaux Y., 2015, ArXiv e-prints, [1502.03120](#)
- Leistedt B., McEwen J.D., Vanderghelynst P., Wiaux Y., 2013, A&A, 558, A128, [1211.1680](#)
- Lewis A., Challinor A., Turok N., 2002, Phys. Rev. D, 65, 2, 023505, [astro-ph/0106536](#)
- Marinucci D., et al., 2008, MNRAS, 383, 539, [arXiv:0707.0844](#)
- Martínez-González E., Diego J.M., Vielva P., Silk J., 2003, MNRAS, 345, 1101, [astro-ph/0302094](#)
- McEwen J.D., Büttner M., Leistedt B., Peiris H.V., Vanderghelynst P., Wiaux Y., 2014, in *IAU Symposium*, volume 306 of *IAU Symposium*, 64–67, [1412.1340](#)
- McEwen J.D., Büttner M., Leistedt B., Peiris H.V., Wiaux Y., 2015, IEEE Signal Processing Letters, 22, 2425, [1508.03101](#)
- McEwen J.D., Durastanti C., Wiaux Y., 2015a, ACHA, submitted, [arXiv:1509.06767](#)
- McEwen J.D., Leistedt B., Büttner M., Peiris H.V., Wiaux Y., 2015b, IEEE TSP, submitted, [arXiv:1509.06749](#)
- McEwen J.D., Vanderghelynst P., Wiaux Y.t.c., 2013, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8858 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 0, [1308.5706](#)
- McEwen J.D., Wiaux Y., 2011, IEEE Transactions on Signal Processing, 59, 5876, [1110.6298](#)
- Moudden Y., Cardoso J.F., Starck J.L., Delabrouille J., 2005, EURASIP Journal on Applied Signal Processing, 2005, 484606, [astro-ph/0407053](#)
- Narcowich F., Petrushev P., Ward J., 2006, SIAM J. Math. Anal., 38, 574, [astro-ph/0403098](#)
- Planck Collaboration, et al., 2014, A&A, 571, A12, [1303.5072](#)
- Planck Collaboration, et al., 2015a, submitted to A&A, [1502.01582](#)
- Planck Collaboration, et al., 2015b, submitted to A&A, [1502.05956](#)
- Planck Collaboration, et al., 2015c, submitted to A&A, [1502.01585](#)
- Planck Collaboration, et al., 2015d, submitted to A&A, [1502.01587](#)
- Planck Collaboration, et al., 2015e, submitted to A&A, [1509.06348](#)
- Planck Collaboration, et al., 2015f, submitted to A&A, [1502.01592](#)
- Planck Collaboration, et al., 2015g, submitted to A&A, [1507.02058](#)
- Rogers K.K., Peiris H.V., Leistedt B., McEwen J.D., Pontzen A., 2016, MNRAS, 460, 3014, [1601.01322](#)
- Seljak U., Zaldarriaga M., 1997, Physical Review Letters, 78, 2054, [astro-ph/9609169](#)
- Smith K.M., Zaldarriaga M., 2007, Phys. Rev. D, 76, 4, 043001, [astro-ph/0610059](#)
- Wiaux Y., McEwen J.D., Vanderghelynst P., Blanc O., 2008, MNRAS, 388, 770, [0712.3519](#)
- Zaldarriaga M., Seljak U., 1998, Phys. Rev. D, 58, 2, 023003, [astro-ph/9803150](#)